An X-Protocol Based Medical Teleconsultation System Using Low or High Speed Networks. A Specific-Design Approach.

Antonio Dueñas, Miguel A. González, Adolfo Muñoz, Carlos H. Salvador.
Lab. Bioingeniería, Clínica Puerta de Hierro, Madrid, Spain.

The objective of this proposal is to provide solutions for the necessities of teleconsultation or telediagnosis among medical professionals, using work stations within the X-Windows environment and applicable in communication lines with an extensive range of bandwidths and operating system independence. Among the advantages sought are savings in transportation, improvement in the quality of the medical attention provided and continued training for the medical professional.

INTRODUCTION

The present state of the technologies dealing with information sciences and telecommunications favors the emergence of proposals aimed at bridging the geographical distances between groups of professionals. In particular, in the case of medical professionals, a typical scenario arises when a doctor reviewing a patient's folder has some doubt about the interpretation of a given document and prefers to consult a colleague with greater expertise in that area. If the specialist is nearby, the doctor can take the document to him. Frequently, however, this is not the case and the only solution is to discuss the problem over the telephone, which does not always clarify the doubt and is often awkward since the document must be described verbally. On these occasions, this model of teleconsultation would be extremely useful. The system described here is part of a larger application, evolution of the one presented in [1], and includes tools for acquisition of documents by scanner, database managing, electronic transfer of patient folders and a user interface designed for medical professionals.

Document

A document will be considered to be any bidimensional image produced by medical equipment. It is digitized by scanning a transparent or opaque original of any size, in color, grays or monochrome. The TIFF file format is employed.

Approach

Systems that deal with the problem of collaboration among professionals that are separated by a physical distance can be classified on the basis of one of two approaches:

1.- Shared windows system: extends existing single user applications by augmenting a window system so that it supports the sharing of the application interface by multiple dispersed users [5].

2.- Collaboration aware application: a specific application designed to directly support multiple cooperating users [3].

The former is more useful for taking advantage of an existing application base for cooperative use. The specific design of the application is considered to be more appropriate for our situation since it allows greater freedom of exact adaptation to the needs of the physicians. These needs are not precisely defined as yet, but will be established throughout the period of evaluation.

On the other hand, from the programming point of view, frequently the development of a communications module focuses too much on the hardware level being too dependent on this or on the operating system, in the attempt to enhance efficacy by taking advantage of the available bandwidth [4]. The solution offered here potentially can be used on any machine with X-Windows, regardless of the operating system, factors which do not imply a significant reduction in performance. The development and evaluation were carried out using a SUN SPARCstation IPX with SunOS 4.1.3.

It is also important to point out that no restrictions or controls have been established with respect to contending user actions, for example, when two users attempt simultaneously to move the document within the window (scroll), one moving it upward and the other downward. These actions are dealt with asynchronously, that is, they are all carried out, but the one that remains is the last one performed by any of the users involved. The proper use of this freedom is left up to the users themselves, who are provided with a voice channel.

DESCRIPTION

To define teleconsultation, it is assumed that it involves a general physician and a specialist or two specialists. The teleconsultation session will develop according to the following steps:

0.- Depending on the speed of the communication line the document(s):
   a) They will have been transmitted previously (if <64 Kb/s).
   b) They will be transmitted immediately before step 1.- (if >64 Kb/s but <2 Mb/s) by means of the utility designed for that purpose.
   c) Each one is sent automatically in step 3 (if >2 Mb/s) (approximate speeds).
1.- Voice contact with the specialist is established by normal telephone, preferably with "free hands" feature to facilitate the use of the mouse. If both agree that to hold the session at that time, then...
2.- The consulting user selects the location corresponding to the specialist from a menu that appears in the document window. The specialist receives notification in the center of his screen, identifying the consulting user and the document. Having confirmed, by means of a button provided in the notification, his acceptance to discuss the document,...
3.- A window, identical to that on the screen of the consulting user, that is, containing the same document in the same position and with the same dimensions, appears on the specialist's screen. From that moment on, any movement of the mouse made over the document by either of the two users will be faithfully reflected in the other screen with a second, larger, colored cursor. The same occurs if the mouse is moved by pressing button 1 (scroll) or 2 (magnifying glass), if the window is moved or iconized, if the brightness of the image is changed, etc. If the cursor goes beyond the edge of the window or if other windows are moved or changed, no modification will appear in the remote screen; thus, the rest of the windows are "private".

The handling by the user of a document or events, which are capable of being reflected in a remote terminal, can be classified according to one of three categories:

1.- Window configuration. This encompasses the events that affect the arrangement of the window on the screen, that is, those that alter the placement or dimension, iconization or visibility.
2.- Directly related to movement of the mouse. The particular features are that:
   a) Temporal peaks with large amounts of information can be produced (when the mouse is moved).
   b) The maximum permitted latency (from the moment the action of the user occurs at one end until it is reflected at the opposite end) must necessarily be short. If it were too long (>0.7 s), it would affect the subjective feeling of agility perceived by the users since there is practically no delay in the voice line. In addition to making it possible to follow the remote cursor by drawing a second cursor in the local window, this category includes the options for scroll, magnifying glass and marking.
3.- Editing of the contents of the window or document as such. This includes those actions that modify the image to improve viewing, for example, a change in brightness or angle. The latency can be somewhat longer than in the preceding category since it is considered together with the time employed in editing, which is greater than in the above cases.

In all three cases, and of special importance in the third, the transmission to the remote station of an order for a change to be made takes place prior to its execution in the local station; thus, the processing time overlaps. For example, the time elapsed between the moment a gamma correction is requested at one end and is carried out at the remote end is only the latency of the communication, not the sum of the latency and the time needed to calculate the gamma correction.

IMPLEMENTATION

The practical use of the above features will be done through the facilities for communications between clients provided by X-Windows. This is possible since, with the client-server philosophy on which X-Windows is based, it makes no difference whether the client and the server are using the same machine or different machines linked by a TCP/IP network since it is the X-Protocol that isolates the system from these differences. The ultimate objective is to take advantage of the suitability of the communications capacities inherent to X11.

There are three possibilities:

1.- **cut-buffers** is a simple but limited form of communication between clients.
2.- **selections** is the most potent and standardized method since release 4 of X [2].
3.- **properties**, used in most cases in communications between the window manager and other clients. Both the cut-buffers and the selections make use of properties.

The direct property system was chosen because it

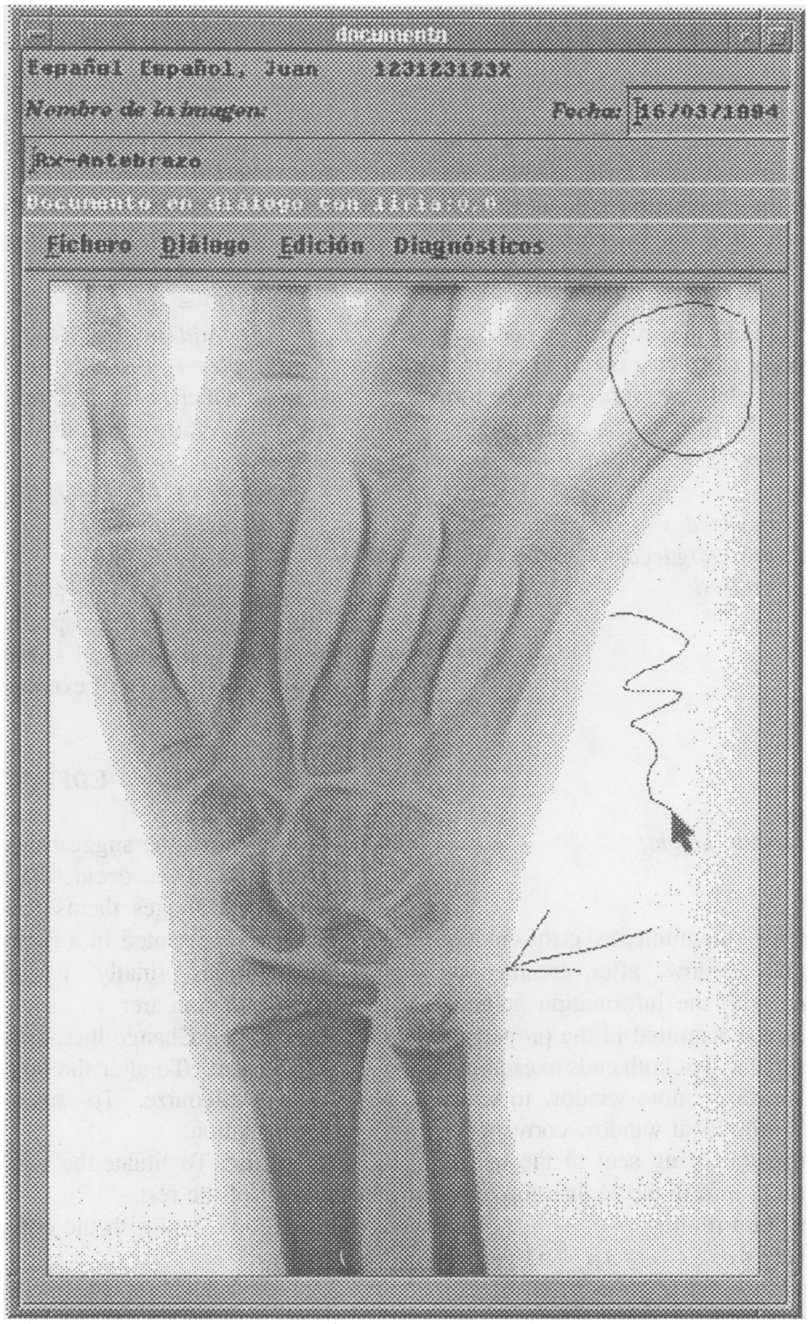

Fig 1: Document window during a teleconsultation session

satisfied our needs since the properties allow the association of arbitrary information with a window, placing it at the disposal of other clients or applications storing it in the server. The properties are referenced by atoms as a means of nickname. A given atom and a given window identify a single property. For properties that are not predefined, in the first call to function **XInternAtom**(), with the name of the property (in the form of a string) as argument, a new atom is assigned to it. This same atom will be assigned to each and every client who calls XInternAtom() with the same name and will remain in the memory of the server while it is active. The properties can contain structures or raw data. They remain set until the window is destroyed, which takes place when the client exits, except in the particular case of a root window which is never destroyed.

Briefly, the procedure would be is follows:

Initially,

myatom = XInternAtom(mydisplay, "TELECONS", *False);*

and whenever an incident is to be communicated to a remote client,

*XChangeProperty(mydisplay, mywindow, myatom, mydeed->type, 8, PropModeReplace, (Deed *)mydeed, sizeof(Deed));*

XFlush(mydisplay);

mywindow will always be the toplevel window of the application, which does not have to be visible. The toplevel window of the application is located in the remote display by searching for the name among all the applications by using **XQueryTree()** and **XFetchName()**.

The **Deed** structure contains the information necessary in each case to properly reproduce the local action in the homologous remote window. It can take the form of a union since it can be different in each case:

```
typedef union _Deed {
    int type;
    AnyDeed myanydeed;
    ConfigDeed myconfigdeed;
    EditDeed myediteed;
    . . .
} Deed;
```

for example:

```
typedef union {
    int type;
    Window w;
    Position x, y;
    Dimension width, height;
    . . .
} ConfigDeed;
```

When the incident being communicated is the creation of a new document window, after creating the homologous window with the information gathered from the data structure transmitted in the property, a table with the two is updated at both ends to establish, from that moment on, the remote window to which the events detected in the local window correspond.

To gather the information being sent in the remote display, it is sufficient to inform the **MainLoop** of the XtToolkit of the attention routine:

XtAddEventHandler(myapptoplevel, PropertyChangeMask, False, deedrcvEH, NULL);

This Event Handler will read the information and proceed to:

*XGetWindowProperty(mydisplay, myapptoplevel, myatom, 0, sizeof(Deed)/4, False, AnyPropertyType, &actual_type, &actual_format, &nitems, &bytes_after, (Deed *)mydeed);*

If this method is applied to the transmission of mouse events, the result is not sufficiently valid for low speed (9600 b/s). As was mentioned above, the events generated by moving the mouse are critical due to their large volume during peak times and the limited tolerance to increasing transmission latency. This is due to the fact that the transmission of properties is of lower priority for the server. Thus, it is necessary to adopt a particular solution for mouse movement which consists in only accepting events occurring within a certain time span for transmission. A value of 100 ms (10 events/s) was empirically selected for this interval, rate which provides more than acceptable subjective quality without exhausting the 9600 b/s limit. To perform all this, a periodical interrupt is installed:

PRDTO = 100;

XtAppAddTimeOut(XtWidgetToApplicationContext(toplevel), PRDTO, perioTO, NULL);

perioTO which, every 100 ms, picks up the last event involving a movement, disregarding those occurring before:

while(XCheckMaskEvent(yourdisplay, PointerMotionMask | ButtonMotionMask, &myxevent));

draw_remote_cursor(): / in local display */*

Thus, regardless of the speed with which the mouse is moved, there will always be a maximum controllable number of events per second, in this case 10.

EDITING

As a result of the suggestions of actual users in the initial trials, it was decided to increase the options for editing the images themselves during the dialogue. These are presented in a menu bar in each window. The utilities finally included for use during teleconsultation are:

1.1.- Move. Change the position of the window.

1.2.- Resize. To alter the height or width.

1.3.- (De)Iconize. To substitute for a symbolic representation.

1.4.- Raise. To situate the window in the foreground, in front of the rest.

1.5.- Close. To finish the consultation regarding this window.

2.0.- Remote cursor. Larger than the local cursor and brightly colored. Always activated.

2.1.- Scroll. If the dimensions of the document surpass those of the window, the window can be moved around over the document by moving the mouse while pressing the left button.

2.2.- Marking. By pressing the middle button, a line is drawn by the cursor as it moves, in both the remote and the local window.

2.3.- Magnifying glass. Pressing the right button augments the area surrounding the cursor to a 2:1

scale.
3.1.- Mirrors. Horizontal and vertical symmetry.
3.2.- Turn. ±90°.
3.3.- Brightness. Changes in intensity.
3.4.- Contrast. Changes in percent or threshold.
3.5.- Gamma correction. $y = x^{**}n$; $n \lesseqgtr \in [0,1]$
3.6.- Restore. Return to the original image.

EVALUATION

The demonstrations and laboratory tests have awakened a great deal of interest among the end users. Several of their suggestions for improving the system have been taken into consideration. However, demonstrating its true validity will require its continuous use within a program of evaluation in circumstances resembling a real situation. At present, the possibilities of the system are being studied in depth in a a pilot demonstrator set up between a health care center and a referral hospital approximately 40 kilometers away. It is tested mainly for transmitting radiologic images (at a rate of 20 a day), problem which is presently being solved by a specialist at the health care center who goes to the hospital twice a week.

CONCLUSIONS

The result is a relatively simple, but robust, application that constitutes the basis for applying teleconsultation in many cases that are currently being handled by less efficient means. Aside from the possibilities of the system itself, there are additional factors, such as:

1.- It is practically independent of the operating system to be used, which means that it can be installed in existing stations. The only requirement is the availability of an X11R4 or higher.
2.- It is also quite independent of the type of communications and its bandwidth. Only TCP/IP or DecNET protocol support is needed, whether it is installed in the same machine or via router in its local network. It has been exhaustively and satisfactorily tested in these cases:
   a) Modem at 9600 or 14400 b/s (V.32 and V.32bis) with error correction (V.42) and via PPP (Point-to-Point Protocol). In this case, the percentage of bandwidth occupancy when the mouses of both users are fully active is nearly 100% (the limit being 10 positions/s and at 9600 b/s), but never reaches saturation.
   b) ISDN at 64 and 128 Kb/s. Equally valid. Connection and disconnection are much more rapid. The surplus bandwidth is serviceable for data and the second B-channel for voice.
   c) Ethernet. In this case, real time transmission of documents is permitted, as well (during the session), avoiding the need to prepare it beforehand. Logically, in other broad band networks (ATM, FDDI), its performance can be expected to be equally satisfactory.

ACKNOWLEDGEMENTS

This work has been supported by the Fondo de Investigaciones Sanitarias de la Seguridad Social, Spain (FIS 92/0412 and FIS 94/0142).

Reference

[1]. C. Blanco, A. Muñoz, A. Dueñas, M.A. Gonzalez, C.H. Salvador. Medical workstation for the management of the transplantation unit of a hospital. proc. Fifth IEEE Symposium on Computer-Based Medical Systems. Durham, 1992.
[2]. D. Rosenthal. Inter-Client Communications Conventions Manual. MIT X Consortium. 1991
[3] J.M. Ng, E. Chang, Horace H.S. Ip, K.Y. Kwok, Y.K. Lee, H.H. Tsang. A Multimedia Conferencing System for Co-operative Medical Diagnosis. Proc. Sixth Annual IEEE Symposium on Computer-Based Medical Systems. 1993.
[4]. S.J. Dwyer. Teleradiology using switched dialup networks. IEEE Journal Selected Areas in Communication (vol. 10, #7). 1992.
[5]. H. Abdel-Wahab, M.A. Feit. XTV: A Framework for Sharing X Window Clients in Remote Synchronous Collaboration. Proc. IEEE Communications for Distributed Applications and Systems. Piscataway, N.J. 1991.